

RADIO NETWORK COMMUNICATION SYSTEM AND
PROTOCOL USING AN AUTOMATIC REPEATER

5

TECHNICAL FIELD

This invention relates to a network of devices which communicate with each other via radio frequency.

10

BACKGROUND OF THE INVENTION

A network of devices can be created by arranging a group of devices that communicate with each other via radio frequency (RF) means to transmit data between the devices. Provided that each of the devices is within the range of the maximum communications range of each device, each device can effectively communicate with each other device in the network.

In many such networks, the devices may all be transceivers, ie are each capable of both transmitting and receiving. For the purpose of transferring an item of data, one

20 device acts predominantly as a transmitter while other devices act predominantly as receivers. In this context, transceivers, which act predominantly as transmitters will be referred to as transceiver/transmitters. Those which act predominantly as receivers will be referred to as transceiver/receivers.

25 The present invention is to be used in a "point to multipoint" system as opposed to a "point to point" system in which communications occur between only two devices at a time. In a point to multipoint communication system, communications occur between one device and two or more of the other devices in the network simultaneously.

30

A reliable "point to multipoint" communication system allows the creation of a shared network variable. This is a variable which is known to all of the devices in

the network. For example, if one device wants to change the value of the shared network variable, it must transmit a request and be guaranteed that all devices receive and process the updated variable simultaneously. If the update cannot be made simultaneously, or not all other devices in the network receive the update, then

5 the network does not have a shared network variable.

Shared network variables allow the creation of a network which has no central controller. All of the essential data about the operation and control of the network is known by each device in the network simultaneously. The data can be updated by

10 any device in the network at any time and all other devices are guaranteed to update their data accordingly. This allows the control of devices within the network to be simplified, more flexible, and less costly when compared with networks having a central controller.

15 A single communication action between each of the devices is herein referred to as a Transaction. A Transaction occurs between a device (transceiver/transmitter) which transmits data to one or more transceiver/receivers of the data. The Transaction also includes data sent from the transceiver/receivers to the transceiver/transmitter as well as to each other of the transceiver/receiver devices in the network.

20 When transmitting to more than one transceiver/receiver simultaneously (also known as a broadcast or multicast), it is important to know that all transceiver/receivers have successfully received the data. If even one transceiver/receiver has not successfully received the data (for example because of a

25 bit error causing data corruption in one transceiver/receiver), then all other transceiver/receivers must be informed that not all of the other transceiver/receivers have successfully received the data.

Such networks may use a transmission system with dominant and inferior bits. This

30 means that if there is a conflict, and two devices simultaneously transmit a dominant

and inferior bit, then when monitoring the communication medium, each device will see the transmission of the dominant bit. The device transmitting the inferior bit knows that there has been a conflict and can take whatever action is appropriate. For example, this may mean the cessation of all further transmissions.

5

In conventional point to point communication protocols, a general procedure is to have each device transmit an acknowledge statement some time after receipt of the data. This has the drawback that the transmitting device must know exactly the number of receiving devices within the network, and how to contact each of them.

- 10 The reliable transfer of the same piece of data to multiple receiving devices requires many transmissions of the same data, and a corresponding wait for each transmission to be acknowledged. The repetitive transmission of the same data to many recipients wastes the available bandwidth of the communication medium. This approach also requires the transmitter to obtain and store data about exactly which
- 15 receiving devices are to accept a given transmission. This approach allows the creation of a shared network variable, at the expense of unnecessary complexity and poor use of the available bandwidth of the communication medium.

Alternatively a point to multipoint transmission can be used to transfer data to many

- 20 recipients simultaneously, without any acknowledgement being returned. This renders the data transfer unreliable, and the transmitter will not be able to determine whether all of the receiving devices have successfully received the data. Unreliable transfer of data means that a shared network variable cannot be created.

- 25 The situation is compounded when two or more devices are out of communications range of each other. As it will be appreciated, each device has a maximum transmitting range (determined by design factors, including but not limited to transmit power, receiver sensitivity, antenna type, and signal processing algorithms). Communications, and synchronising communications is made more complicated

when some devices are outside the maximum range and therefore cannot communicate with each other.

It is an object of the present invention to provide a system and protocol for

5 improving the communications between devices in an RF multicast communications system, particularly between devices which are disposed beyond the normal communications range of one or more of the other devices in the system.

SUMMARY OF THE INVENTION

10 According to a first aspect of the present invention, there is provided in a radio communication system including a first transceiver, a second transceiver and a repeater, the first and second transceivers being separated from each other by a distance greater than at least one of their respective maximum transmission ranges, and the repeater being located intermediate the first and second transceivers, the

15 method including upon receiving data from one of either the first or second transceivers, the repeater transmits a repeater flag and then transmits the data received from the one of either the first or second transceivers.

According to a second aspect of the present invention, there is provided a

20 communications protocol for use in a network of devices, the protocol having a frame including a first time slot for transmitting data, a second time slot, after the first time slot, for indicating a repeat flag, and a third time slot, after the second time slot, for retransmitting the data transmitted in the first time slot.

25 According to a third aspect of the present invention, there is provided a radio communication system including a first transceiver, a second transceiver and a repeater, the first and second transceivers being separated from each other by a distance greater than at least one of their respective maximum transmission ranges, and the repeater being located intermediate the first and second transceivers,

30 wherein upon receiving data from one of either the first or second transceivers, in a

first time slot, the repeater transmits a repeater flag in a second time slot, and then in a third time slot transmits the data received in the first time slot.

According to a fourth aspect of the present invention, there is provided a repeater for
5 use in a radio communication system including at least two transceivers, the at least
two transceivers being separated from each other by a distance greater than at least
one of their respective transmitting ranges, in use, the repeater being disposed
intermediate the at least two transceivers wherein upon receiving data in a first time
slot, the repeater transmits a repeat flag in a second time slot, and then transmits in a
10 third time slot, data received in the first time slot.

According to a fifth aspect of the present invention, there is provided a transceiver
for use in a radio communication system including at least one other transceiver and
a repeater, the transceiver and the at least one other transceiver being separated from
15 each other by a distance greater than at least one of their respective transmitting
ranges, in use, the repeater being disposed intermediate the transceiver and the at
least one other transceiver, wherein upon receiving a repeat flag from the repeater, in
a second time slot, the transceiver suspends further action until it receives from the
repeater, in a third time slot, data that was originally transmitted by the at least one
20 other transceiver in a first time slot, before the second time slot.

According to a sixth aspect of the present invention, there is provided a method for
use in a radio communications system including at least a first transceiver, a second
transceiver and a repeater, the first transceiver and the second transceiver being
25 separated by a distance greater than a maximum transmission range of at least one of
the transceivers, and the repeater being disposed intermediate the first
and second transceivers, such that upon receipt of a data transmission from the first
transceiver, the repeater re-transmits the data transmission from the first transceiver,
wherein upon receipt of a data transmission from the second transceiver before the
30 repeater completely receives or retransmits the data transmission from the first

transceiver, the repeater transmits a data sequence instructing each transceiver to cease its respective transmission.

According to a seventh aspect of the present invention, there is provided a radio

5 communications system including at least a first transceiver, a second transceiver and a repeater, the first transceiver and the second transceiver being separated by a distance greater than a maximum transmission range of at least one of the transceivers, and the repeater being disposed intermediate the first and second transceivers, such that upon receipt of a data transmission from the first transceiver,

10 the repeater re-transmits the data transmission from the first transceiver, wherein upon receipt of a data transmission from the second transceiver before the repeater completely receives or re-transmits the data transmission from the first transceiver, the repeater transmits a data sequence instructing each transceiver to cease its respective transmission.

15 According to an eighth aspect of the present invention, there is provided a repeater for use in a radio communication system including at least a first transceiver and a second transceiver, the first transceiver and the second transceiver being separated by a distance greater than a maximum transmission range of at least one of the transceivers, in use, the repeater being disposed intermediate the first and second transceivers such that upon receipt the data transmission from the first transceiver, the repeater retransmits the data transmission from the first transceiver, wherein upon receipt of a data transmission from the second transceiver before the repeater completely receives or re-transmits the data transmission from the first transceiver,

20 the repeater transmits a data sequence instructing each transceiver to cease its respective transmission.

25 According to a ninth aspect of the present invention, there is provided a transceiver for use in a radio communication system including at least one other transceiver and a repeater, in use, the transceiver and the at least one other transceiver are separated

by a distance greater than a maximum transmission range of at least one of the transceivers, and the repeater being disposed intermediate the transceiver and the at least one other transceiver, such that upon receipt of a data transmission from the at least one other transceiver, the repeater re-transmits the data transmission from the
5 at least one other transmitter and upon receipt of the data transmission from the transceiver before retransmitting the data transmission from the at least one other transceiver, the reporter transmits a data sequence instructing each transceiver to cease respective transmissions, wherein, upon receipt of the data sequence from the repeater, the transceiver will cease transmission.

10

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 shows one network architecture according to a preferred embodiment of the present invention;

Figure 2 shows the network protocol model used in the environment of the present

15 invention;

Figure 3 shows a frame structure according to a preferred embodiment of the present invention;

Figure 4 shows a repeater and two transceivers disposed in a network according to the present invention;

20 Figure 5 shows a frame structure used in the configuration of Figure 4;

Figure 6 shows a preferred structure of a transceiver device used in the present invention; and

Figure 7 shows parts of the ISO 7 layer model in which certain functions of the transceiver device of Figure 6 are performed.

25

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

An exemplary architecture of a network is shown in Figure 1, in which the network
10 is made up from nodes 1, 2 and 3. Nodes 1, 2 and 3 are transceiving devices and may act as transmitters and/or receivers in a given communication Transaction. The

30 network 10 may communicate with other networks 20, via gateway 11.

The protocol design of the present invention is based on the ISO seven layer model and some terminology is common with that used by ISO. The protocol used in the present invention is connectionless, meaning that once a single data transfer has

5 taken place, there is no expectation of additional related data transfers either before or after.

The protocol model of the present invention is based on the ISO seven layer model and is shown in Figure 2. For applications distributed over two nodes, each protocol

10 layer has a virtual connection to the equivalent layer in the other node. As can be seen, each layer takes data provided by the layer above, treats it as a data unit and adds its own protocol control information (PCI) field. At each layer, the protocol data unit (PDU) is either the data, or a package provided by the next higher layer.

15 The name of the PDU is prefixed by the layer to which it applies (For example, SPDPU is a session PDU).

The physical layer relates to the mechanical and electrical network interface in the ISO system. In the system of the present invention, the physical layer refers to the hardware and firmware elements used to transmit and receive bits over the

20 communication medium.

In the ISO system, the link layer is used for data link control (for example, framing, data transparency, error control). In the present invention, the link layer is used to break bytes up into bits, bit stuffing (if required), framing, collision detection,

25 prioritisation, error detection, positive/negative acknowledge generation, checking, repeating and retransmission.

The network layer in the ISO system is used for networking routing, addressing, call set-up and clearing while in the present invention, the network layer is used for

30 network routing, addressing, Transaction set-up and clear.

In the ISO system, the transport layer is used for end to end message transfer, connection management, error control, fragmentation and flow control. The transport layer is not used in the environment of the present invention.

5

The session layer in the ISO system is used for dialogue and synchronisation control for application entities but is not used in the environment of the present invention.

The presentation layer is used for transfer syntax negotiation, and data

10 representation transformations in the ISO system while in the environment of the present invention, the presentation layer is used for optional encryption of application data.

The application layer in the ISO system is used for file transfer, access management,

15 document and message interchange, job transfer and manipulation while in the environment of the present invention, the application layer supports sending and receiving application data.

Finally, the user application layer is used both in the ISO and the environment of the

20 present invention for whatever is needed to achieve a specified function or behaviour.

It is predominantly in the link layer that the features of the present invention reside.

25 In the protocol of the present invention, use can optionally be made of a dominant bit, and an inferior bit. If two devices simultaneously transmit a dominant and an inferior bit, then receivers and transmitters (monitoring their own transmissions) will detect only the dominant bit.

Media access is obtained by a transmitter first monitoring the media, and if no existing transmission is detected, the transmitter will try to claim media access by transmitting a preamble stream. This preamble starts with at least one detectable bit.

5 The claim for media access defines the start of a Transaction. A Transaction consists of all data transfer, acknowledgement and repeating of data. All nodes in a network must monitor the media continually and if they detect a Transaction occurring they will defer any attempt to claim media access until the completion of the current Transaction.

10 Transactions are asynchronous: they can occur at any time and the time difference from the start of one Transaction to the next does not have to be an integral number of bit periods.

In this application, a Transaction is specifically defined as a continuous period of

15 time broken up into several sub-time slots containing different types of data. A Transaction will begin with a preamble for a set period of time, followed by the specific data which is to be transmitted from a transceiver/transmitter to two or more transceiver/receivers. The timeslot during which the data is transmitted is variable in length, and includes a portion used as a frame check sequence. Following
20 the data transmission is a timeslot during which positive acknowledgement can be transmitted by the transceiver/receivers, followed by a timeslot during which negative acknowledgement is transmitted by the transceiver/receivers. The structure of this frame is shown in Figure 3.

25 As described above, a Transaction is asynchronous and can start at any time. However, once started, the Transaction has a time-based structure. Special markers in the Transaction are used to show the beginning and end of the variable length data portion. The time slots during which the positive and negative acknowledgement are transmitted, are fixed in time. By coding and redundancy of data encoded into
30 these timeslots, a positive acknowledgement by one or more transceiver/receivers

and a negative acknowledgement by one or more transceiver/receivers can be conveyed. All of the devices involved in the Transaction see both of the acknowledgement timeslots.

- 5 Transceiver/receivers wishing to positively acknowledge, will transmit a special code during the positive acknowledge timeslot and during the negative acknowledge timeslot will either receive (if dominant bit transmission is not used) or transmit inferior bits (if dominant / inferior bits are used).
- 10 Similarly, transceiver/receivers wishing to negatively acknowledge will either transmit inferior bits (if dominant / inferior bits are used) or receive (if dominant bit transmission is not used) during the positive acknowledge timeslot and transmit a special code during the negative acknowledge timeslot.
- 15 The fact that the devices monitor the timeslots they are not transmitting, means that by the end of two acknowledge timeslots each device has detected either positive acknowledges, negative acknowledges or both and can therefore work-out the overall acknowledge state of the network.
- 20 For example, the transceiver/receiver which transmits a positive acknowledge will be able to detect some other transceiver/receiver which transmits a negative acknowledge. For the case where dominant bits are used, the transceiver/receiver positive acknowledging will attempt to transmit inferior bits during the negative acknowledge timeslot but will detect dominant bits due to the other
- 25 transceiver/receiver which is simultaneously transmitting an appropriate code during the negative acknowledge timeslot. For the case where dominant bits are not used, the transceiver/receiver positive acknowledging will receive during the negative acknowledge timeslot and detect any bits due to the other transceiver/receiver which is transmitting an appropriate code during the negative

acknowledge timeslot. The reverse case applies for the transceiver/receiver transmitting a negative acknowledge.

At the end of the Transaction, all devices do not know how many positive or

5 negative acknowledges there were, all they have to know is that there were some positive and some negative.

If there were any negative acknowledges at all during the Transaction, then all of the transceiver/receivers know this, and can discard the data received. Similarly, the

10 transceiver/transmitter knows this and can attempt to re-run the Transaction.

The generation of a positive acknowledge will be as follows. Upon receiving data, a node will generate a positive acknowledge only when:

-the data timeslot has been checked against its embedded frame check sequence and found to be valid; and

15 -any addressing information present in the data timeslot matches an addressing information used by the device.

Each device transceiver generally contains at least two different types of address, as follows:

20 - A Unit Address, allowing the device to be uniquely addressed in isolation; and

- A Multicast address, allowing those devices in a network to be addressed simultaneously for the purpose of updating shared network variables.

In addition, devices can also optionally contain:

25 - A Network Address, allowing physical devices to be grouped by the logical network to which they are allocated.

Other variations are possible, but these three address types are used as the basis for other more sophisticated addressing schemes.

The processes involved in generation of a negative acknowledge are as follows. A receiving device (transceiver/receiver) will generate a negative acknowledge only if the data timeslot is determined to be corrupted, by checking the received data using the embedded frame check sequence.

5

Where a device determines data timeslot corruption, there is no point further examining any fields inside the data timeslot.

Acceptance of the data transmitted by the transceiver/transmitter is only made by

10 the transceiver/receivers if the condition for generation of a positive acknowledge are met and if no other transceiver/receiver has generated a negative acknowledge. This ensures that all transceiver/receivers receive a given message only once. For point-multipoint messages, this may mean that a message is discarded by a transceiver/receiver, even if it appears valid and was positively acknowledged.

15

The above describes the general environment in which the present invention can be utilised. The sequences described above can only be used when each device is within range of the other. It will be understood that each device will have a maximum transmitting range, beyond which it cannot communicate with other devices. The

20 maximum transmitting range is determined by design factors, including but not limited to transmit power, receiver sensitivity, antenna type, and signal processing algorithms. For short range (unlicensed) devices, the range is typically from tens to, at most, several hundred metres. In the case where one or more devices are located beyond the maximum transmitting range of another device (ie cannot communicate

25 with that device directly), difficulties in implementing the above-described procedure will be encountered. Particularly, in the case where a transceiver/transmitter transmits data, some or all of the other transceivers/receivers will not receive the data from that particular transceiver/transmitter and accordingly, will make it impossible to update shared
30 network variables.

In accordance with an aspect of the present invention, the protocol described above (and subject of a co-pending patent application) is modified to allow the retransmission of data between devices so as to extend the effective transmission

5 range of the devices used in the network. The modified protocol is used in conjunction with a repeater which is placed roughly in the geometric centre of devices in the network and acts as a relay between devices distributed beyond their normal transmission range.

10 Figure 4 shows an example configuration of devices A and B in a network. Devices A and B are separated by a distance greater than each of their respective transmitting ranges. Thus, if device A were to transmit data as described above, device B would not receive this data and would not know how to proceed as described above. This would prevent the effective updating of a shared network variable. In accordance

15 with an aspect of the present invention however, repeater device 40 is placed between devices A and B and acts as a repeater. Thus, if device A transmits data, repeater device 40 will receive this transmission from device A and retransmit the data such that device B will receive device A's data. When device B transmits its acknowledgement, this will be received by the repeater. The repeater in turn

20 transmits an overall acknowledge status which will be received by both devices A and B. Both devices then know that the information was relayed by the repeater, and in turn accepted or rejected by all devices in range of the repeater. Devices A and B can then proceed in the normal manner.

25 Of course, device B for example need not be a transceiver/receiver but may be a transceiver/transmitter. In this case, device B will transmit information to the network however, because device A (for example a transceiver/receiver) being out of range of device B would not receive the transmitted data. Again, repeater device 40, being disposed between device A and device B, will receive the data transmitted by

30 device B, and retransmit this data so that device A and any other devices within the

range of repeater device 40 will receive the retransmission. Similarly, the acknowledgement from device A will be received by the repeater, which in turn provides an overall acknowledge status back to device B.

5 It will be appreciated that in practice, repeater device 40 does not need to be placed directly between two devices but may be placed in any suitable position such that devices within the network can be reached.

In some cases, it may be possible that device A will have sufficient range to reach

10 device B however, device B, having a shorter transmission range than device A, will not be able to communicate with device A. In this case, repeater device 40 may be positioned closer to device B than to device A in order to allow transmissions from device B to reach repeater device 40 which can then be repeated and communicated to device A.

15

Practically, it is beneficial to construct all devices in the network in the same way.

This means that each device, whether it acts as a transceiver/transmitter, transceiver/receiver or repeat device, will be constructed in the same way and can be separately enabled to perform their desired functions. This provides significant

20 savings in the complexity and cost of manufacture since only one device need be manufactured. The particular construction of a device will be described in more detail later, with reference to Figures 6 and 7.

In use, if a device has been set as a repeater, upon receiving information in the first

25 frame (see Figure 3), the repeater will immediately transmit a repeat flag in a new second time slot and then retransmit, in a new third time slot, the data received in the first time slot. The network then operates normally as described above in that devices being transceivers/receivers that have received the retransmitted information will then proceed to acknowledge the successful or unsuccessful receipt
30 of that data as discussed above, and the repeater will issue a final overall repeat

status to inform all devices in the network of the success or failure of the repeated data.

The modified protocol frame is shown in Figure 5. A comparison with the frame of

5 Figure 3 clearly shows the differences between the modified repeat tag transaction of Figure 5 and the non-repeat tag transaction of Figure 3. Specifically, the first time slot providing for transmission of data is present in both transaction frames however, in the repeat tag frame of Figure 5, a second time slot is provided for the transmission of a repeat tag flag. A third time slot is provided in which the data

10 transmitted in the first time slot is then retransmitted. An acknowledgement time slot is then provided containing a first sub-time slot for transmitting a positive acknowledge and a second sub-time slot for transmitting a negative acknowledge. Furthermore, in the repeat tag frame of Figure 5, an additional time frame is provided for the transmission of a repeat status which provides confirmation to all

15 devices that the transmission has been repeated.

The situation described above in relation to the repeater is complicated even further by the fact that it is possible that some devices will begin the transmission at the same time. When all devices are within range of each other, collisions can normally

20 be dealt with easily when a system of dominant and inferior bits is used. Firstly, collisions are avoided by monitoring the transmission medium before transmitting. This leaves a small period during which simultaneous transmission can begin. In the event that two devices simultaneously begin transmission in this small period, there will inevitably be a difference in the data bits being transmitted. When this difference

25 occurs, the device that transmits the inferior bit will detect a dominant bit due to the other transmitter, and can then cease further transmission.

The device that detects the collision reschedules its transmission for some later time. This time delay can be based on a small random number, optionally scaled by the

30 length of the message to be transmitted.

In the case where a repeat device is used, there is a much longer delay between a device transmitting and another device receiving that transmission via the repeater.

- 5 For example, with reference again to Figure 4, if device A begins transmitting, there will be a delay between the time repeater device 40 receives device A's transmission and the time it retransmits the transmission to be received by device B. During this time, device B may begin transmitting its own data. This new transmission will cause a collision, which is detectable at the repeater but not at device A. This situation is
- 10 addressed by another aspect of the present invention.

Specifically repeat device 40 checks the data that it receives. If it detects a bit stuffing violation or a data coding violation during its reception, this indicates that devices A and B are transmitting simultaneously. Upon detection of this, repeater device 40

- 15 deliberately begins transmitting a long stream of dominant bits (for example 6 to 8 bits), which violate the normal bit stuffing or data coding rules. This will cause devices A and B to both detect a collision, and cease transmission according to normal collision detection and resolution rules as described above.

- 20 As for a normal collision, Devices A and B then each schedule a retransmission after a random time delay, optionally scaled according to the message length. The likelihood of this delay being identical for each device is very small however, should each device again begin transmitting at the same time, the delay period is recalculated, but this time, each device multiplies its respective delay period by two.
- 25 For each subsequent failure, the successive delay periods are doubled up to some predetermined number of doublings. At this point, the transmission will be aborted and optionally the operator can be notified by any suitable means. In most cases however, the random delay period will result in a resolution to the conflicting transmissions and the devices of the network will be able to continue transmitting in
- 30 the normal manner.

Of course, it will be appreciated that any other suitable form of re-transmission delay could be used.

- 5 As previously discussed, practically, it is beneficial to construct all devices in the network in the same way. This means that each device, whether it acts as a transceiver/transmitter, transceiver/receiver or repeat device, will be constructed in the same way and can be separately enabled to perform their desired functions. This provides significant savings in the complexity and cost of manufacture since only
- 10 one type of device need be manufactured.

The preferred implementation of transceiver device 100 uses a radio receiver, a radio transmitter, and a microprocessor. These first two items can optionally be combined as a transmitter/receiver, as shown in Fig 6, which shows device 100 including

- 15 microprocessor 110 and transmitter/receiver 120. Transmitter/receiver 120 transmits and receives data via antenna 130.

It will be appreciated that use of a microprocessor is not mandatory. For example the protocol could be implemented in a dedicated integrated circuit, a programmable

- 20 logic device or a programmable gate array. Using a microprocessor is convenient because it allows a readily modifiable software implementation, and reduces the overall parts count. However the software implementation is only suitable for low to moderate data rates.

- 25 The function of the transmitter/receiver 120 is to receive or transmit information. The choice of transmitter/receiver will be determined by a range of factors, including (but not limited to):
 - a. The regulatory environment of the market into which the product will be sold.

Each country has regulations that determine factors including allowable frequencies, transmit power levels and bandwidth.

Transmitter/receivers suitable for use in some countries may be illegal in other countries.

5 For products with a broad sales appeal in a range of countries, it is possible that several different transmitter/receivers will need to be selected, appropriate to each country.

- b. Power consumption, in conjunction with any other considerations that determine the amount of power available.

10 For example, transmitter/receivers with a high power consumption may be unsuitable for battery operation.

- c. Time for the transmitter/receiver to switch between receive and transmit modes.

In the communication protocol of the present invention, the time to switch between receive and transmit is important, because the protocol includes a set of fixed time slices. Depending on the overall transaction being performed, a time slice may need to be received or transmitted.

15 The time to switching between receive and transmit constitutes an overhead (dead time). Large switching times have the effect of wasting bandwidth.

- d. Interface type.

Many types of transmitter/receiver are available. A type that provides digital data input and output gives the simplest interface with the microprocessor.

- e. Data rate.

The transmitter/receiver needs to support a data rate appropriate to the overall product requirements. This data rate could be anywhere between extremely low or extremely high.

- f. Physical size, and the amount of space available.
- g. Cost.
- h. Amount of design effort.

5 At the very least, the transmitter/receiver needs:

- a. A Transmit data output, used by the microprocessor to place a communication state onto the wireless medium;
- b. A Receive data output, used by the transceiver to indicate to the microprocessor the state of the wireless medium; and
- c. A control input, used by the microprocessor to select a receive or transmit mode of operation of the transmitter/receiver.

Control inputs can range between very simple, and very complex. At the simplest extreme, it is used to select between receive and transmit. Some transmitter/receivers
15 support a low power "sleep" mode. Others allow complex setup and configuration to be made for transmitter/receiver operating behaviour.

For the protocol described, the type of control input is not critical.

20 Some suitable transmitter/receivers include RFM ASH series TR1000 to TR3100, Chipcon CC1000 and Nordic NRF401, NRF403.

The microprocessor 110 is used to implement the communication protocol, using the transmitter/receiver as the means of placing communication states onto the wireless
25 medium, and receiving communication states from the wireless medium.

The type and choice of microprocessor is not very critical, provided that it can perform operations with precise timing. The degree of precision only needs to be enough to avoid creating bit errors in the communication protocol.

5 The protocol is best implemented in a bit-oriented manner, because this allows the point at which the time slices begin to be easily recognised.

The microprocessor is responsible for implementing at least some of the following functions:

10 a. A data encoding and decoding scheme used for transmission and reception - for example, Manchester coding;

b. Recovery in the receiver of the transmitted clock - for example by synchronising onto a preamble;

c. Detection of collisions;

15 d. Creation of each time slot, and appropriate transmission or reception during the time slice to exchange the relevant acknowledge information;

e. Implementation of an error detection scheme that can be used by a receiving device to determine if a transmission is received with or without error;

20 f. Implementation of an error correction scheme that can be used by a receiving device to correct for some number of received errors during a transmission; and

g. Adding a repeater function that can be used to change the structure of the transaction, allowing the information packet to be repeated for the purpose of extending the range.

As previously described, a common method for describing the functions used in communication protocols is the ISO 7 layer model. Whilst a software structure based on this model is not mandatory, its use simplifies the overall design. Using this model, the functions performed in the bottom few layers are shown in Fig 7.

5

The microprocessor hardware provides the electrical interface (Physical layer), while the microprocessor software performs the functions of all higher layers.

In particular, the software MAC-B portion of the Link Layer is responsible for all

10 time-critical functions of data transmission and reception, including at least some of the following:

- a. Starting a new transmission (including generation of any preamble);
- b. Transmitting data bits;
- c. Transmitting frame markers;
- d. Starting reception;
- e. Synchronisation onto the transmitted data stream and clock recovery;
- f. Reception and decoding of data bits;
- g. Reception and decoding of frame markers;
- h. Detection of a collision;
- i. Starting each time slice; and
- j. Transmission or reception of data bits inside a time slice.

The software MAC-F portion of the Link Layer is not so time critical. It is responsible

25 for higher level message-oriented processing, including at least some of:

- a. Construction of a packet from received data bits;

- b. Packet error checking;
 - c. Determination of when to acknowledge, and the type of acknowledgement to be generated (using the time slice services of MAC-B);
- 5 d. (optionally) Based on packet structure, scheduling the operation of a repeated transmission
- e. Starting the transmission of a new packet;
- f. Generation of packet error-check sequences;
- g. Transmission of a packet, bit at a time;
- 10 h. Examination of acknowledgements, and collisions, and determination of whether a packet should be re-transmitted and when that re-transmission should occur.

Many different microprocessors are available. Some have special hardware functions available which remove some of the processor load for time critical functions, such as generation of time intervals, generation of pulses, and so on. Whilst these hardware functions are not mandatory, their use greatly simplifies the software design and coding.

20 Some microprocessors that are suitable for use in this invention include the Texas Instruments MSP430 family, Atmel Atmega family and the Hitachi H8/3644 family.

A useful feature of an aspect of the present invention relates to the accurate detection of the end of the variable length time slice containing the data transmitted by the transceiver/transmitter.

It is desirable for communication media to require some form of balanced transmission to avoid accumulation of a dc offset. This balance requires that the

number of ON and OFF states on the medium be equal, when considered over a moderate to long time period.

5 There are a number of encoding schemes that can be used to translate the data bits into states on the medium. These vary in terms of the bandwidth they consume on the medium, and the ease of recovery of the transmitted data in a receiver.

One of the most common is Manchester coding. This coding uses two states on the medium for each data bit, and has a simple process for data recovery in the receiver.

10 This codes a 1 bit as the state pair (OFF, ON), and a 0 bit as the state pair (ON, OFF).

The Manchester code always has a state transition (OFF to ON, or ON to OFF) in the middle of each data bit which greatly simplifies the process in the receiver of data recovery and synchronisation onto the transmitter clock.

15

In the Manchester code, the state pairs (OFF, OFF and ON, ON) are not permitted.

The Manchester code can be exploited, so that the illegal state pairs are used to convey information about points of significance.

20

The exact choice of the illegal state sequence is not very important, so long as it is used consistently. Preferably, the dc balance of the Manchester coding should be preserved.

25 A suitable coding to represent the end of the variable portion of the transmission is to use a simple illegal sequence: (ON, ON, OFF, OFF). This preserves the dc balance, and can be easily recognised by the Manchester decoder.

If additional information needs to be conveyed, this sequence can be used as a "lead-

30 in". So, for example, other possible sequences might be:

(ON, ON, OFF, OFF, ON, OFF) = first point of significance

(ON, ON, OFF, OFF, OFF, ON) = second point of significance

5 When the methods and advantages are considered together, the preferred protocol implementation is bit-oriented, synchronous, and exploits illegal coding to denote points of significance in the variable part.

This has the advantage of providing a high level of time-based precision in finding

10 the end of the variable portion, is relatively easy to implement, and does not need escape sequences or bit-stuffing. Furthermore, the high time-based precision in finding the end of the variable portion also creates a high level of precision in determining the start of the fixed time slots that follow.

15 The fixed time slices are easily transmitted, simply by counting the transmitted states or bits. The received time slices require a Manchester decoder (without the need to support the illegal states), and a timer in the case where nothing is transmitted during the period of the received time slice.

20 It will be understood that the above has been described with reference to a preferred embodiment and that many variations and modifications may be made within the scope of the present invention.